UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/635,703	08/05/2003	David B. Glasco	NWISP036	8389

22434 7590 10/05/2005

BEYER WEAVER & THOMAS LLP
P.O. BOX 70250
OAKLAND, CA 94612-0250

EXAMINER	
THOMAS, SHANE M	

ART UNIT	PAPER NUMBER
2186	

DATE MAILED: 10/05/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/635,703	GLASCO, DAVID B.	
	Examiner	Art Unit	
	Shane M. Thomas	2186	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *02 September 2005*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-20* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-20* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date *8/1/2005*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Continued Examination Under 37 CFR 1.1 1 4

A request for continued examination under 37 CFR 1.1 14, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.1 14, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 3 7 CFR 1.114. Applicant's submission filed on 8/1/2005 has been entered.

Claims 1-20 remain pending; claims 1,2,4,11, and 20 have been amended.

All previously outstanding objections and rejections to the Applicant's disclosure and claims not contained in this Action have been respectfully withdrawn by the Examiner hereto.

Information Disclosure Statement

The information disclosure statement filed 8/1/2005 has not been fully considered by the Examiner because of the following inconsistency: U.S. Patent No. 6,108,797 is not directed to Sharma et al. but to Lin et al. and due to the ambiguity, has not been considered by the Examiner. However, based on the Office Action mailed May 12, 2005, from related U.S. Application No. 10/288,347, submitted on 8/1/2005 and listed on the present IDS, page 51 of the submission shows a copy of an IDS filed with the '347 Office action and lists the Sharma reference with correct patent number (6,108,797). The Examiner has considered the Sharma reference and listed it on the form PTO-892 enclosed herewith.

Response to Amendment

As per Applicant's amendment to the claims, the Bauman et al. reference (U.S. Patent No. 6,189,078) has been cited to teach a --completion indictor-- bit (claim that is used when a response can be handled in a local cluster. Further, the Sharma reference has been cited to reject claims 1-7,11-17, and 20, and combined to with Keller to reject claims 8-10,18, and 19.

Response to Arguments

Applicant's arguments with respect to claims 1-20 have been considered but are moot in view of the new grounds of rejection.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

Claims 1-7, 11-17, and 20, rejected under 35 U.S.C. 102(b) as being anticipated by Sharma et al. (U.S. Patent No. 6,108,737).

As per claim 1, with reference to figure 1, Sharma teaches receiving a cache request associated with a memory line (cache line) at a cache coherence controller 180 from a requesting processor (102-108) from a cluster of processors that are connected by switch 200 (i.e. point-to-point architecture) in column 7, line 60 - column 8, line 8. Response information for the cache request from a remote data cache (i.e. one of the other private caches 122-128 of the processor

cluster) that is associated with the cache coherence controller (since they are all part of the same processing cluster as shown in figure 1) is obtained as taught in column 8, lines 3-8.

The cache coherence controller can determine if the cache access request can be handled locally by examining the address of the memory line requested in order to determine if the memory line is locally owned or remotely owned (i.e. by one of the other processor clusters (SMP Node) as shown in figure 4 (refer to column 10, lines 44-65). Remote nodes (i.e. a cache coherence controller from another SMP Node) will not have to be probed if the home cache coherence controller determines that the request can be satisfied locally (column 10, line 53 - column 11, line 7).

A completion indicator (commit-signal - figures 3 and 8) is sent with response information when the cache access request can be handled locally (column 9, lines 15-36).

As per claim 2, Sharma teaches that the cache access can be handled locally if a valid copy (e.g. dirty copy) of the memory line is in the remote data cache (one of the private caches 122-128 that does not belong to the requesting processor). Refer to column 15, lines 34-59.

As per claims 3 and 13, response information can include state information. For example, when a request to modify data is received by the cache coherence controller from a requesting processor (CID command), a response is a CTD-Success (column 7, lines 35-43). Such a response indicates that the state of the requested data is "exclusive" and can be modified by the requesting processor. Thus, the response of CTD-Success includes information indicating that no other processor cache currently contains valid copy of the memory line.

As per claim 4, the completion indicator (commit-signal) can be a completion bit as shown in figure 3, element 310.

As per claim 5 and 15, the completion indictor (if a type 0 Commit-signal) notifies the processor that the response from the cache coherence controller will be the only response since with a type 0 commit-signal, no external probes (i.e. global invalidates or fetch requests) are outstanding since the request is a local request.

As per claims 6 and 16, the Examiner is considering the processor making the memory line request to be a --requesting processor-- and the cluster (SMP node) that comprises the requesting cluster to be a --requesting cluster--. Refer to column 15, line 60 - column 16, line 8.

As per claims 7 and 17, the completion indicator allows the cache coherence controller to avoid probing local nodes (i.e. processors) since the shared memory 150 of a local cluster can own the requested memory line if no other processor in the cluster does not contain a valid copy of the data (column 6, liners 33-38). Further, if the request is a local request, targeted to the distributed memory portion of the home node (column 10, lines 44-65), the cache coherence controller does not have to probe remote nodes (via a global request) to another one of the SMP nodes (figure 4).

As per claims 11 and 20, the rejections follows the combination of the rejections for claims 1 and claim 2 (wherein the cache coherence controller provides response information to a requesting processor it if it is determined that a valid copy of the memory line is in a remote data cache - as defined supra).

As per claim 12, the response information can comprise a response packet (column 5, lines 57-59).

As per claim 14, Sharma teaches that the response information can include data (column 7, lines 35-43).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 1-7, 11-17, and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hum et al. (U.S. Patent Application Publication No. 2004/0123047) in view of Bauman et al. (U.S. Patent No. 6,189,078).

As per claims 1, 11, and 20, Hum teaches a method for providing a response to a cache access request in figures 4 and 6. Further, Hum shows a *processing cluster* 110 in figure 1. Figure 4 shows the receiving of a cache access request [that is associated with a memory line] at a cache coherence controller (agent 120, figure 1). The cache coherence controller 120 represents multiple nodes, and cluster interface 210 (figure 2) provides two-way communication between the agent 120 and the set of nodes (¶76). Since Hum states that an agent represents multiple nodes (¶76), that a node can be comprised of a single processor (¶109), and that the cache coherency system of the present invention is scalable (¶71), the Examiner is considering elements 112-118 to be simple nodes (comprised of a single processor). Thus *cluster* 110 comprises cache coherence controller 120 and processor nodes 112-118. The processors of cluster 110 are interconnected in a point-to-point architecture through agent 120 (refer to ¶¶ 55, 57 and 76). Response information for the cache access can be obtained from remote data cache 250 (import data cache) in accordance to figure 4 or can be obtained from remote data cache 260

(export directory) in accordance to figure 6. Since figure 2 is a representation of an agent (cache coherence controller) is can be seen that remote data caches 250 and 260 are associated with a given agent. In accordance with figure 4, the agent provides response information (step 450) to the requesting node (processor).

A determination regarding whether the cache request can be handled locally by using a remote data cache 250 or 260 without having to probe remote nodes is taught by Hum in ¶67.

Hum does not specifically teach in ¶67 that if the request is handled locally, that a completion response is provided with the response information (claim lines 10-11). Completion responses are well-known in the cache-coherency art to enable a requesting processor determine whether it must wait [for invalidating acknowledgements from other processors caching the data] in order to modify the requested data, thereby preserving data coherency among the processing cluster. Bauman teaches such a method when response information is sent from a cache coherency controller, taking the form of two message types: a Type 1 response (column 17, Table 1) and a Type 3 M_RESPONSE (column 19, Table 2). Both responses contain a --completion response-- (PG_ACK bit - column 18, lines 1-53), which is used to indicate (when valid) to the requesting processor that all necessary purges (invalidations from other processor's caches that cache the requested data) have occurred or (if set in the type 1 response) that no purges are necessary. Such purging (invalidating) functionality is in line with the Hum reference in that requests are sent to the clusters that contain the data that a requesting processor is requesting access (¶¶64-67). Bauman teaches that the two messages, when used together, minimize latencies associated with delivering the requested data to the requesting processor (column 19, lines 40-44).

Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to have combined the cache coherency system of Hum with the completion response messaging system of Bauman in order to have minimized the latencies in delivering requested data to requesting processors. Such a benefit readily applies to the scenario regarding a processor requesting data for read modification purposes (i.e. a request that can be fulfilled by the export cache [remote data cache 260]). It could have been seen that such a request would have received a response from the agent (cache coherency controller) with the PG_ACK bit [of modified Hum] set if no other Peer node contained the data, thereby indicating that no delayed purge will occur and the requesting processor can immediately modify the data. Without the completion response (PG_ACK bit) indicating that no other processors currently cache a valid copy of the data, the requesting processor would have to wait for NACK messages to be returned from each of the Peer Nodes (¶66 of Hum), thereby slowing down the time until the requesting processor can utilize the data.

The messaging of modified Hum that contains the completion response also is similar to the ACK message of Hum (¶44). The ACK message is taught to indicate that the requested data has been sent to the requesting node, whereas the Type 1 response is also used to inform a Node (denoted --POD-- in Bauman) that requested data is being transferred to the requesting Node in response to the fetch request. (column 17, lines 51-53).

Further regarding claims 11 and 20, the cache access request can be handled locally if a valid copy of the memory line is in the remote data cache as taught in figure 4 (steps 420-450) and figure 6 (steps 620-640) of Hum.

As per claim 2, the cache access request can be handled locally if a valid copy of the memory line is in the remote data cache as taught in figure 4 (steps 420-450) and figure 6 (steps 620-640) of Hum.

As per claims 3 and 13, Hum teaches in ¶43 that in the cache coherency system of the present invention, state information (M/E/F/S) can be included as part of the response information.

As per claims 4, as shown in Tables 1 and 2 of Bauman, the --completion indicator-- is a --completion bit-- (PG_ACK bit), which designates whether all necessary purges have been completed.

As per claims 5 and 15, if the --completion indictor-- is set (PG_ACK = 1) in the response from the modified agent of Hum, it would indicate that the current response will be the only response from the agent (cache coherency controller) since the delayed purge acknowledge signal 912 will not need to be sent to indicate that all delayed purges have occurred (column 18, line 1 - column 19, line 1.

As per claims 6 and 16, the Examiner is considering the node (processor) requesting the data to be a --requesting processor-- (for instance processor 112, figure 2), and the requesting processor's cluster (cluster 110) to be a --requesting cluster-- since the cluster comprised the processor from which a request for data initiated.

As per claims 7 and 17, as shown in figure 4, if the requested data is in the import cache (remote data cache) of the cache coherence controller, the data is forwarded to the requesting node (step 450). Only if the request misses in the remote data cache will the cache coherence controller (120) send the request to the nodes that are represented by the cluster (step 445).

Further, the --completion indicator-- allows the cache coherence controller to avoid probing local or remote nodes when a data request (without write privileges) for data that is cached in the remote cache (250 or 206) occurs, as the cache coherency controller maintains the status of the requested data in local and remote nodes (¶66). The completion indicator is shown being set for a read request in column 18, lines 46-48.

As per claim 12, the Examiner is considering the data response message from the agent 120 (cache coherence controller) that is sent to a requesting node (processor) to be a --response packet-- that comprises the requested data and the state of the data (M/E/F/S) to be used by the requesting processor (¶43).

As per claim 14, Hum teaches that the response can include data (¶43).

Claims 8-10, 18 and 19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Sharma et al. (U.S. Patent No. 6,108,737), as applied to claims 1-7, 11-17, and 20 above, in further view of Keller (U.S. Patent No. 6,728,841).

As per claims 8 and 18, Sharma not specifically teach the [requesting] processor sending a --source done-- upon identifying the completion indicator in the response from the cache coherency controller. The source done response would have allowed the target node to remove the request for data from its command queue and can proceed to process the next request for the memory location. As evidenced by Keller, the sending of the acknowledgement (source done) from the source node back to the target node allows a cache coherent multiprocessor system to process memory read requests effectively. Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to have combined the cache response

method of Sharma with the messaging response teaching of Keller in order to have effectively maintained cache-coherence during memory reads in the multiprocessing system of Sharma (figure 1 for instance). It would have been seen by one having ordinary skill in the art that once the requesting node of modified Sharma (source node of Keller) received the read response (Fill - column 7, line 38) containing requested data followed by an type 1 commit-signal (i.e. completion signal - column 15, line 60 - column 16, line 8) from a remote data cache associated with a remote coherence controller, the requesting processor (node) would have sent a --source done-- response back to the remote cache coherence controller so that the remote cache coherence controller could have removed the read request from its command queue and to proceed to grant the subsequent read request.

As per claims 9 and 19, as discussed in the rejection for claims 8 and 18, the --source done-- message is sent to the remote cache coherence controller since it is the --ordering point-- for incoming cache requests and for maintaining cache coherency for the memory portion of the distributed memory space for which it resides (column 6, lines 58-67, and column 19, lines 4-14).

As per claim 10, as discussed in the rejection for claims 8 and 18, the --source done-- message is sent to the remote cache coherence controller since it is the --ordering point-- for incoming cache requests and for maintaining cache coherency for the memory portion of the distributed memory space for which it resides (column 6, lines 58-67, and column 19, lines 4-14). The Examiner is considering the cache coherence controller of each SMP node to be --acting as a memory controller-- since the cache coherency controller is responsible for interrogating the DTAG 160 [to determine which processor of the system owns the cache line

and which entities have copies of the line] and for controlling the forwarding all memory responses to the requesting processors (refer to column 7, line 60 - column 8, line 8).

Claims 8-10, 18 and 19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Hum et al. (U.S. Patent Application Publication No. 2004/0123047) in view of Bauman et al. (U.S. Patent No. 6,189,078), as applied to claims 1-7, 11-17, and 20, above, in further view of Keller (U.S. Patent No. 6,728,841).

As per claims 8 and 18, modified Hum does not specifically teach the [requesting] processor sending a --source done-- upon identifying the completion indicator in the response from the cache coherency controller. Keller teaches in column 2, line 63 - column 3, line 11, that a step of sending a --source done-- upon receiving a response packet (including requested data) from a target node. The source done response would have allowed the target node to remove the request for data from its command queue and can proceed to process the next request for the memory location. As evidenced by Keller, the sending of the acknowledgement (source done) from the source node back to the target node allows a cache coherent multiprocessor system to process memory read requests effectively. Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to have combined the cache response method of Hum with the messaging response teaching of Keller in order to have effectively maintained cache-coherence during memory reads in the multiprocessing system of Hum (figure 1 for instance). It would have been seen by one having ordinary skill in the art that once the requesting node of modified Hum (source node of Keller) received the read response (containing requested data) followed by an type 1 message containing the PG_ACK bit (completion

indicator) from the remote data cache (import cache 250) of the cache coherence controller (agent 120), the requesting processor (node) would have sent a --source done-- response back to the cache coherence controller so that the cache coherence controller could have removed the read request from its command queue and to proceed to grant the subsequent read request.

As per claims 9 and 19, as discussed in the rejection for claims 8 and 18, the --source done-- message is sent to the cache coherence controller (agent 120).

As per claim 10, as discussed in the rejection for claims 8 and 18, the --source done-- response is sent back to the cache coherence controller (agent 120). The Examiner is considering the cache coherence controller to be --acting as a memory controller-- since it controls the read requests send to a node in its local cluster and can forward the data back to the requesting node without having to send the request to the target node (processor). Refer to ¶63 of Hum.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shane M. Thomas whose telephone number is (571) 272-4188. The examiner can normally be reached M-F 8:30 - 5:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt M. Kim can be reached on (571) 272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Shane M. Thomas

HONG CHONG KIM
PRIMARY EXAMINER